Industry Trends in Microprocessor Design

H. Peter Hofstee

IBM Cell/B.E. Chief Scientist

October 4, 2007

Outline

- Why Hybrid
- Cell/B.E. & roadmap
- Industry direction towards C/GPU & Hybrid systems
- Conclusions

Why Hybrid

SPECINT: Slowdown in single thread performance growth

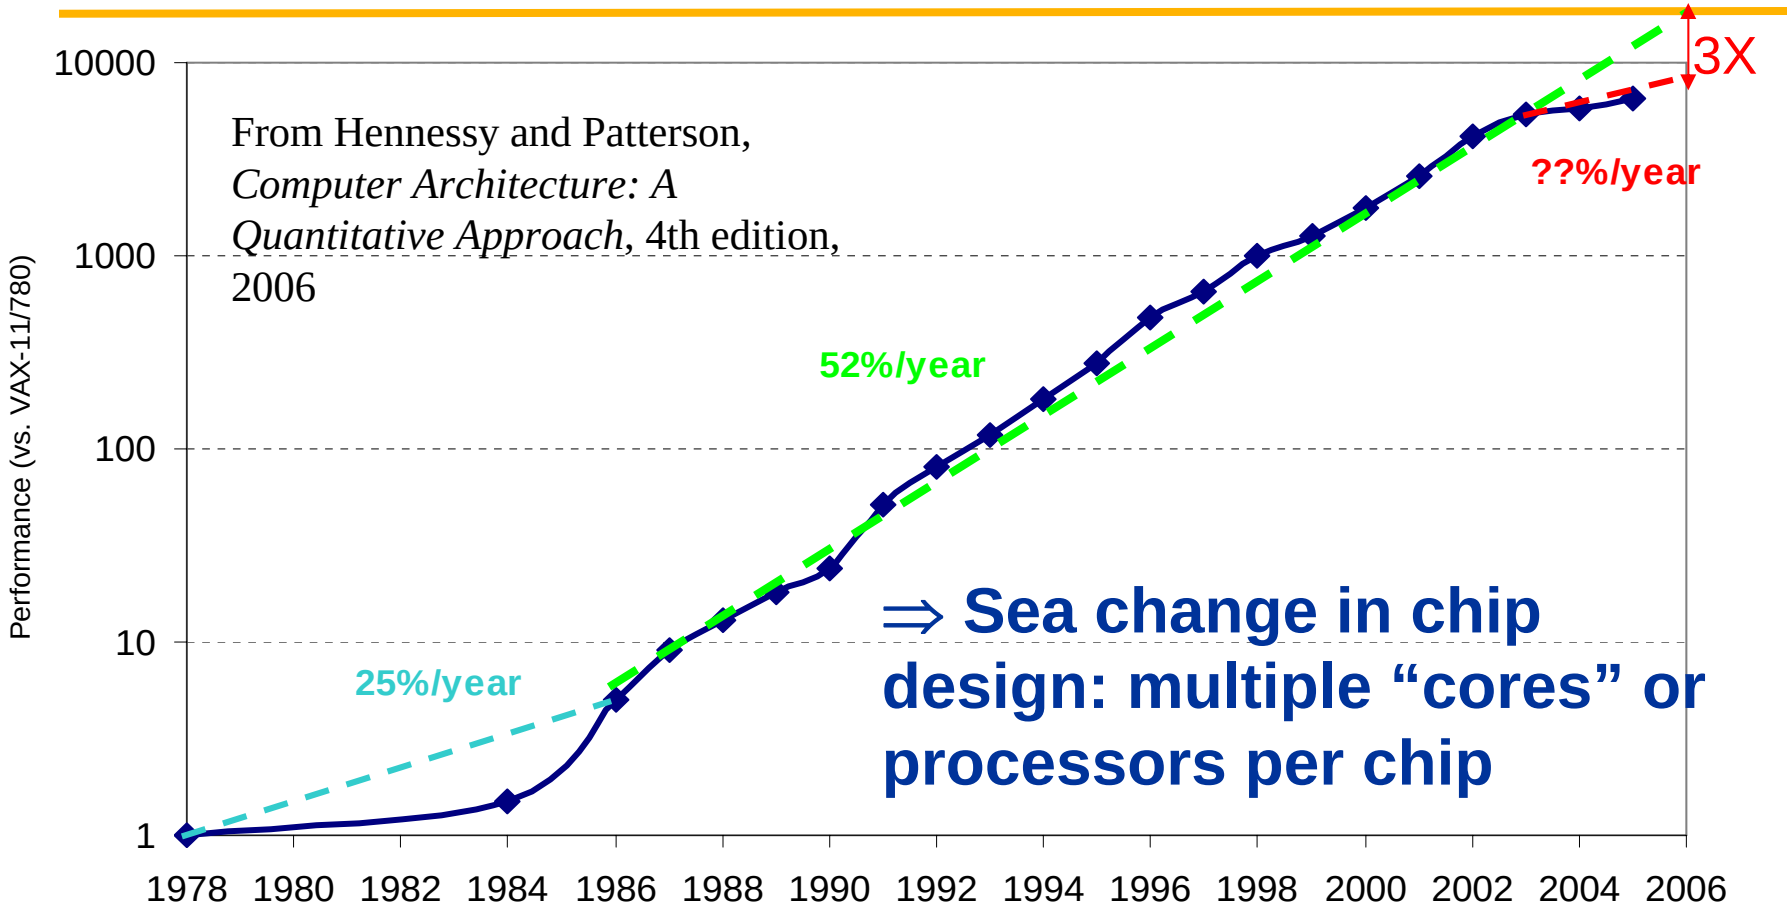

Microprocessor Trends

- Single Thread performance is power limited
- Multi-core extends throughput performance
- Hybrid extends both performance and efficiency

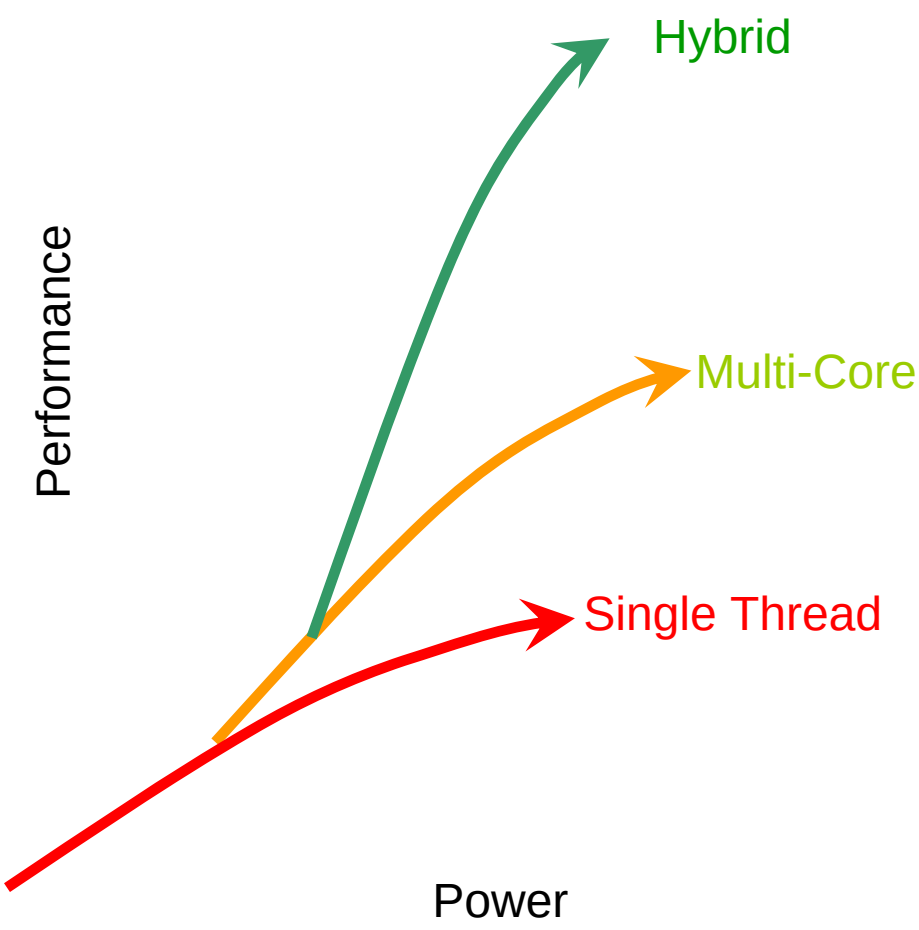

Traditional General Purpose Processor

IBM Power5+

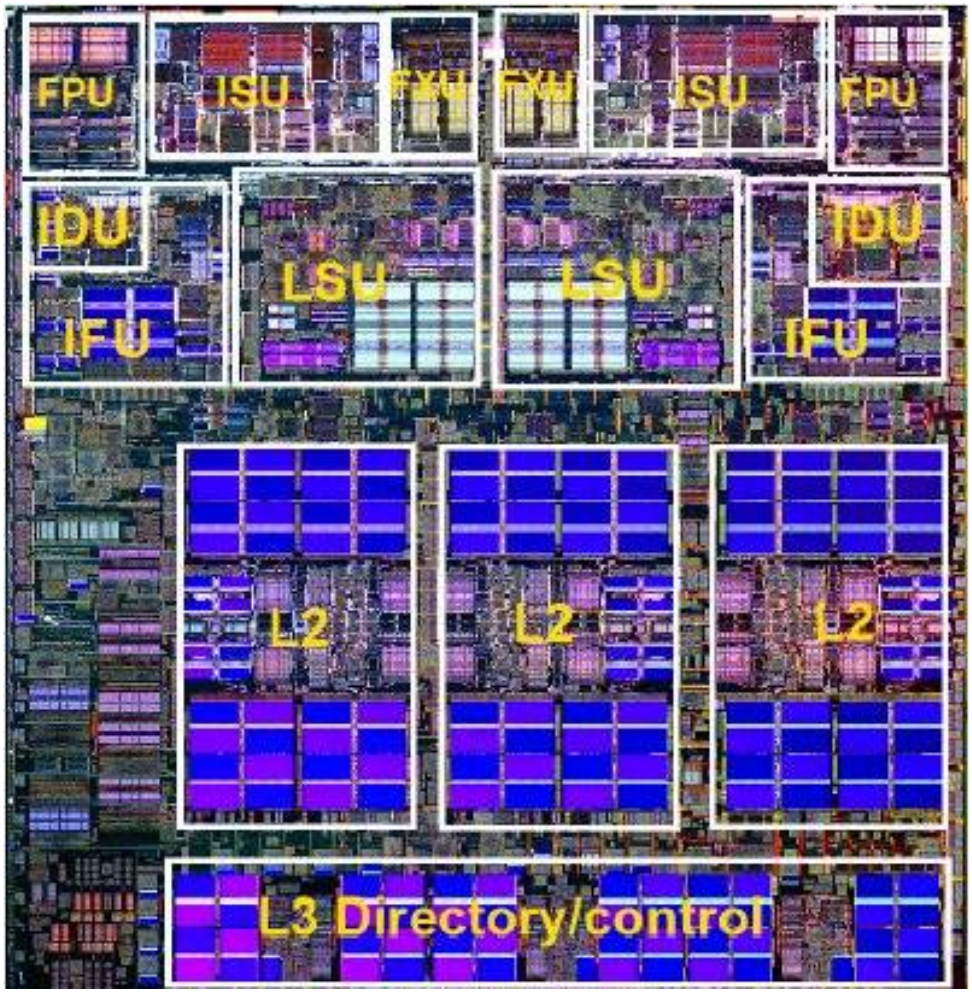

Memory Managing Processor vs. Traditional General Purpose Processor

Cell/ B.E.

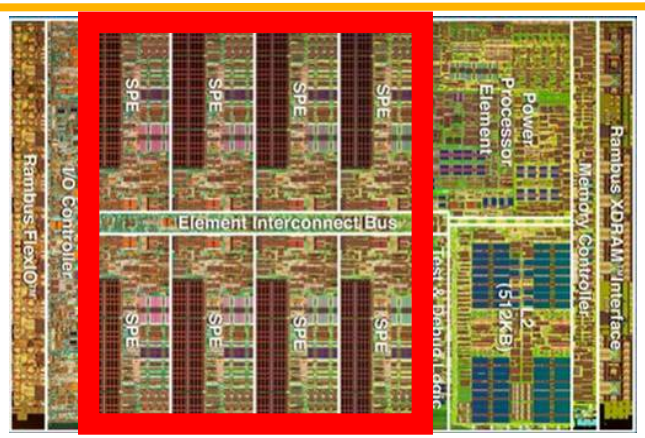

AMD

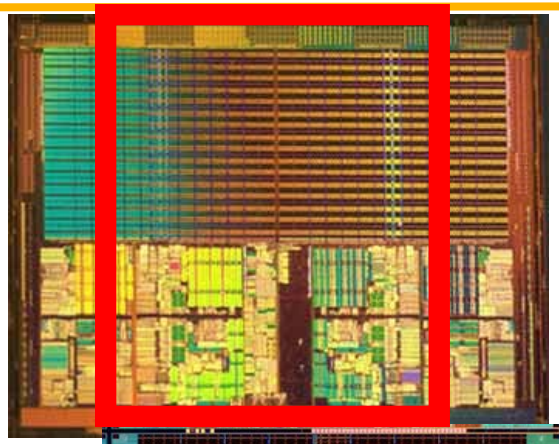

IBM

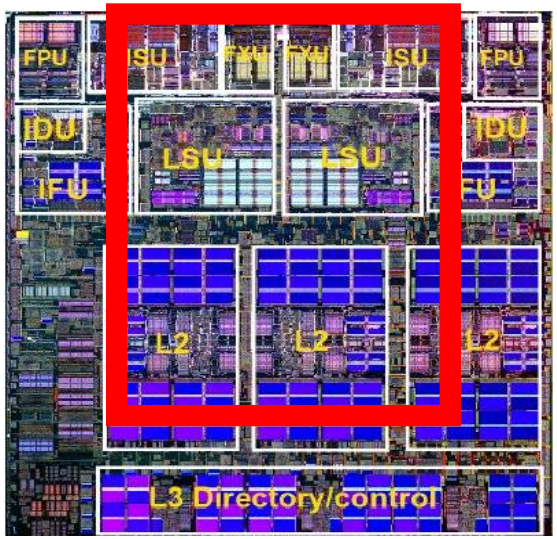

Intel

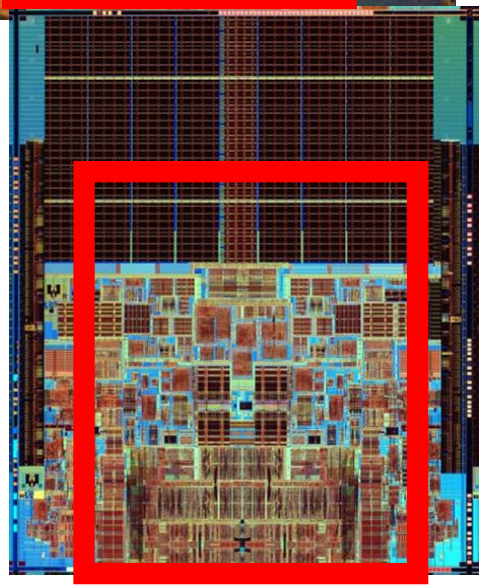

Cell/B.E. & Cell Roadmap

SPE Highlights

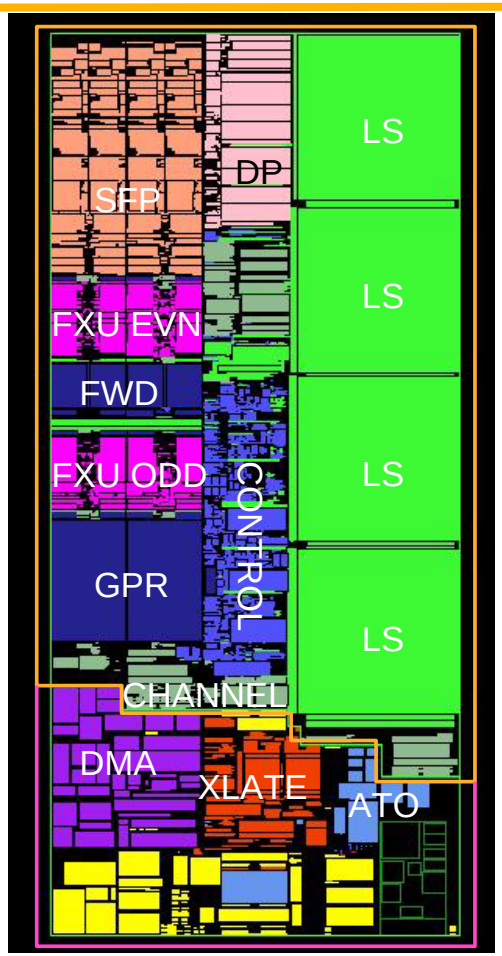

14.5mm² (90nm SOI)

- Flexible DMA Engine
  - Improve effective memory bandwidth
  - Vector Load/Store w/ Scatter Gather
  - DMA is full Power Arch protect/x-late
- 256 KB Local Store
- Not just a coprocessor, has its own PC
  - RISC like organization
  - 32 bit fixed width instructions
  - Dual Issue, high design frequency design
  - Broad set of operations (8/16/32/64)
  - VMX-like SIMD dataflow
  - DP-Float support
- Large unified register file
  - 128 entry x 128 bit (I&FP)
  - Deep unrolling to cover unit latencies
- User-mode architecture
  - No need to run the O/S
  - No translation/protection within SPU

Cell Broadband Engine ™: A Heterogeneous Multi-core Architecture

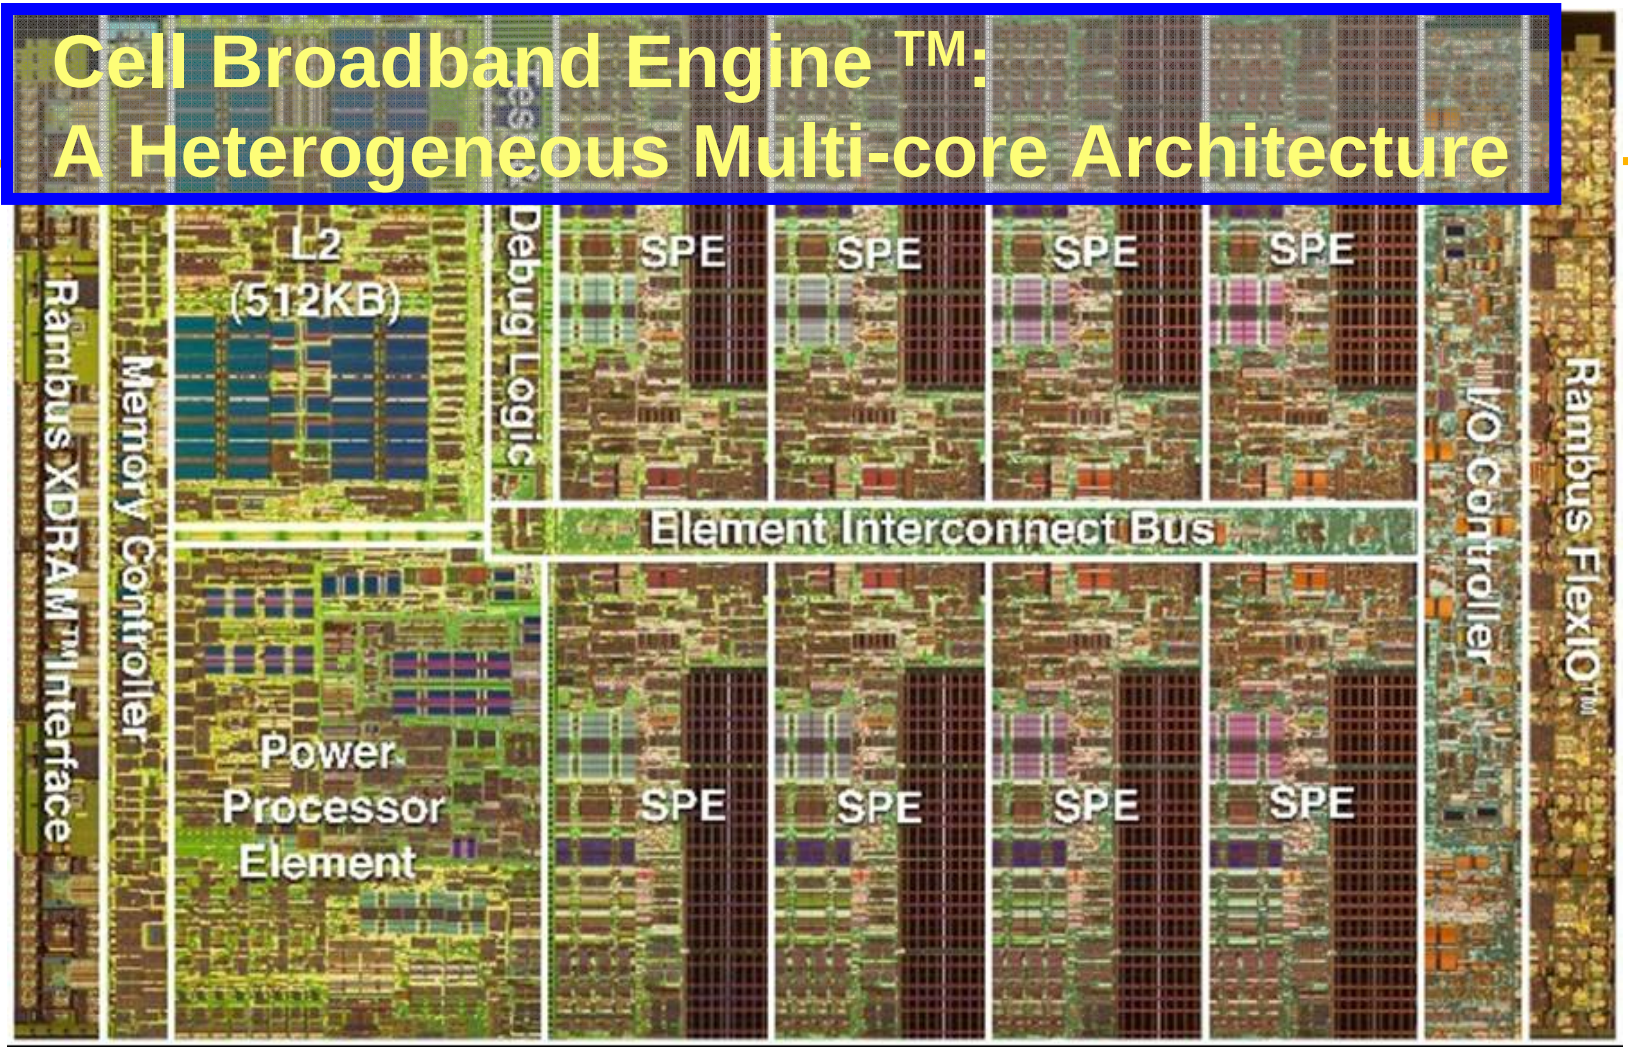

** Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc.*

Cell Broadband Engine™ Architecture (CBEA) Technology Competitive Roadmap

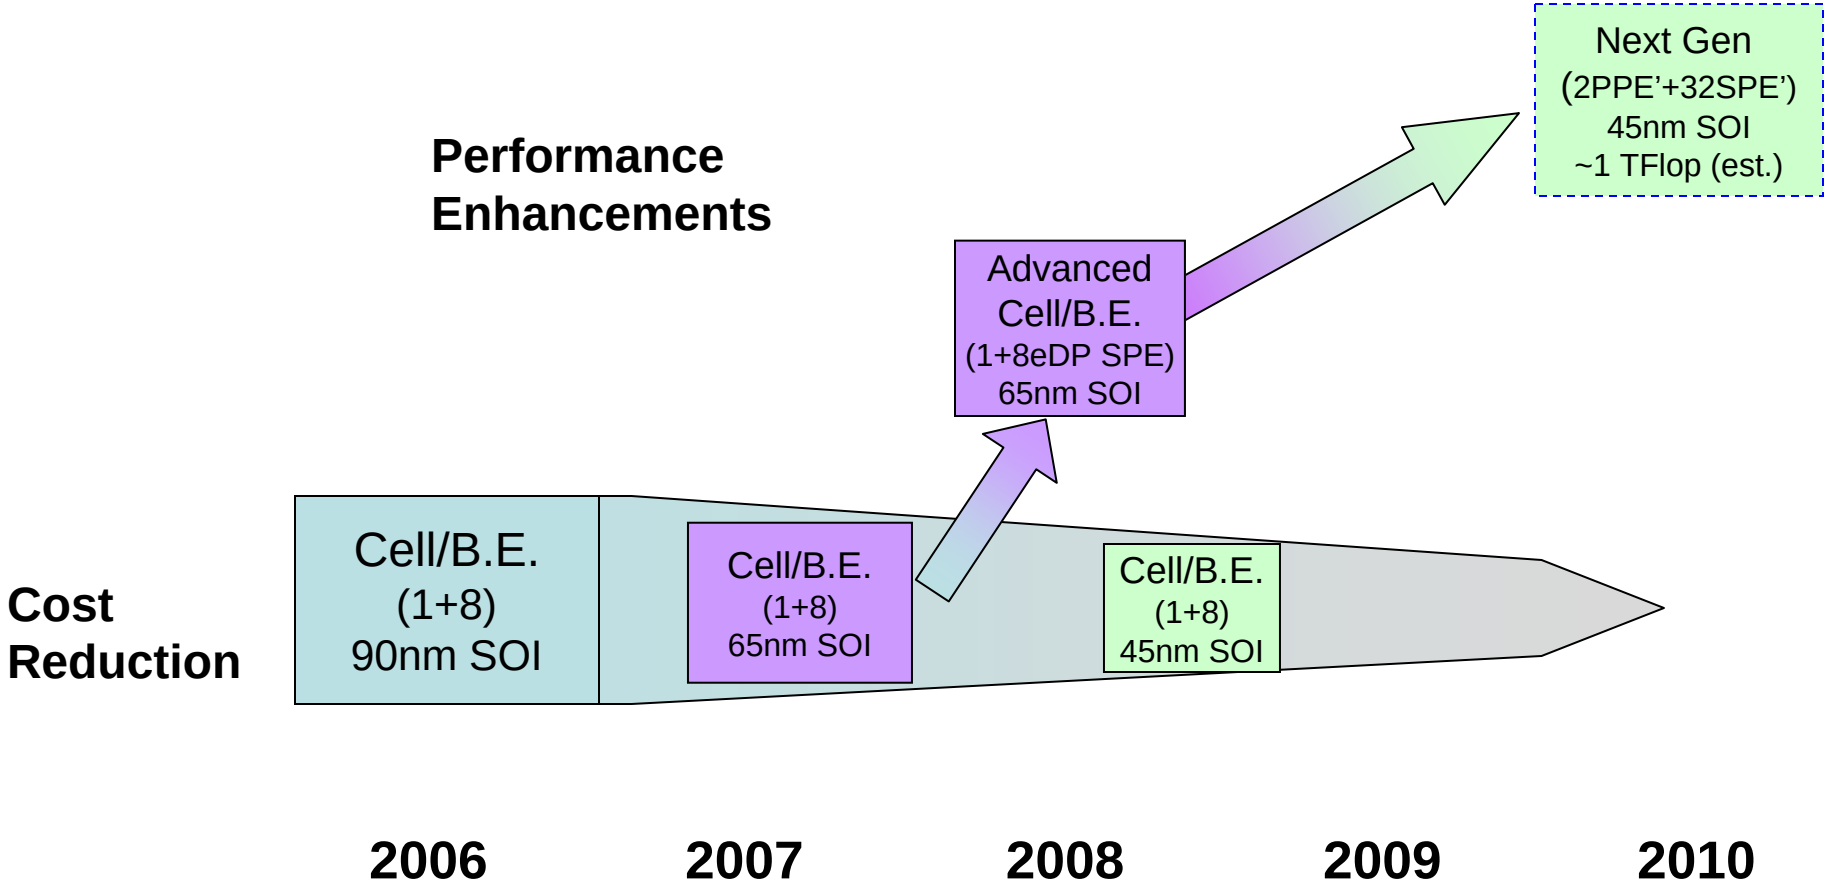

All future dates and specifications are estimations only; Subject to change without notice. Dashed outlines indicate concept designs.

Enhanced BE (DDR2 & an SPE with enhanced DP-FLoat)

Challenge	*Response*
2GB Memory Limit	DDR2 allows upto 16 GB
Still want 25 GB/s	Upto DDR2-800 Many more pins
25.6 Gflops DP/BE •13 Cycle DP Latency •6 Cycle Stall •No Dual Issue w/DP	102 Gflops DP/BE •9 Cycle DP Latency •Fully Pipelined DP •Dual Issue w/DP
IEEE compliance •Denormal Inputs -> 0 •Default NaNs	IEEE compliance is improved •Denormal Support •Expected NaNs
Up to 10% perf. In DP Compare Emulation	5 new DP compare instructions – SPU ISA v1.2

Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

MICROPROCESSOR REPORT ANALYSTS' CHOICE

Best High-Performance Embedded Processor

STI Cell BE

Cell Processor Isn't Just for Games.

Innovative Chip is best high-performance embedded processor of 2005

We chose the Cell BE as the best high-performance embedded processor of 2005 because of its innovative design and future potential....Even if the Cell BE accumulates no more design wins, the PlayStation 3 could drive sales to nearly 100 million units over the likely five-year lifespan of the console. That would make the Cell BE one of the most successful microprocessors in history.

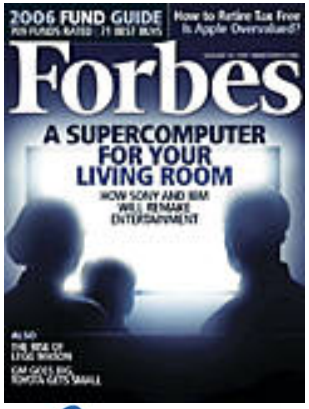

"…Cell could power hundreds of new apps, create a new video-processing industry and fuel a multibillion-dollar build out of tech hardware over ten years."
-- Forbes

"It was originally conceived as the microprocessor to power Sony's [PS3], but it is expected to find a home in lots of other broadband-connected consumer items and in servers too."
-- IEEE Spectrum

Los Alamos NATIONAL LABORATORY EST.1943

IBM

Cell/B.E. based Systems: SCEI, IBM, Mercury, …

Los Alamos NATIONAL LABORATORY EST.1943

IBM ASC NNSA

QS20

QS21

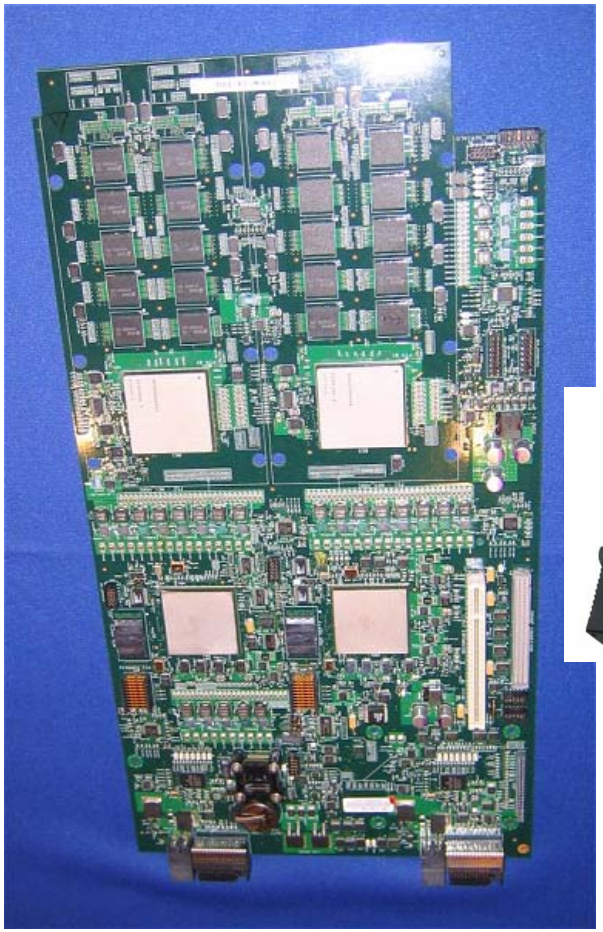

Many Applications for Cell/B.E. Beyond Gaming

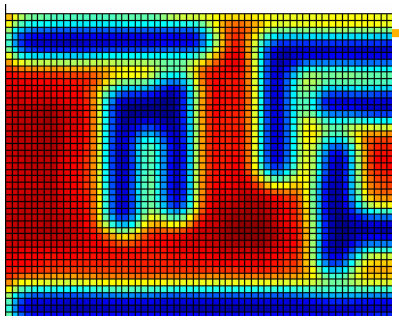

Mercury/Mentor Graphics
45nm OPC tool

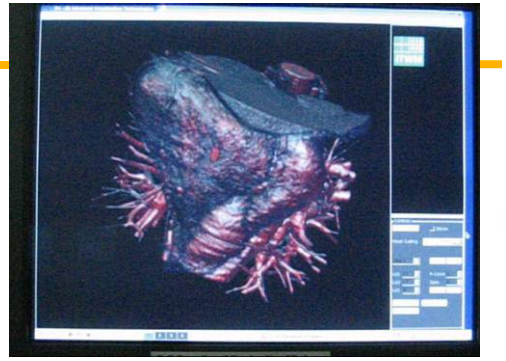

Fraunhofer
PV4D Medical Imaging

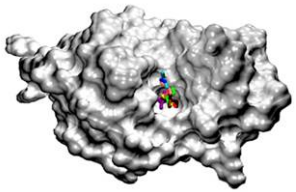

Boston Univ.
Bioinformatics: FBDD

I.B.M. to Build Supercomputer Powered by Video Game Chips
By JOHN MARKOFF
(NY Times): September 7, 2006

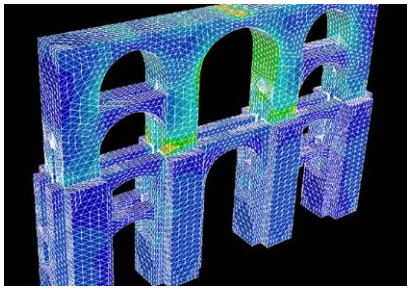

Structural Analysis
digitalmedics.de

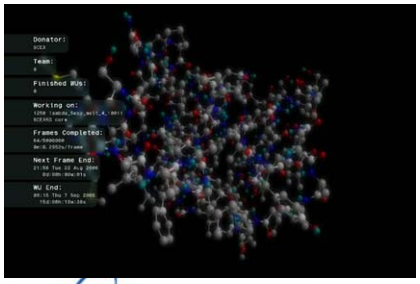

SGE / Pande (Stanford)
folding@home PS3 client

Rapidmind(TM) / RTT

IBM iRT raytracer prototype

Industry workloads well-suited to Cell/B.E. technology

EDA

Financial markets

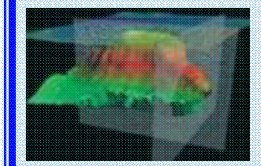

Seismic

A&D

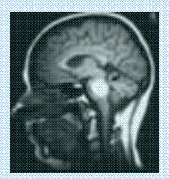

Medical Imaging

Digital Video Surveillance

Media & Entertainment

Focused Common Workload Characteristics/Requirements

Real-time Analytics
Processing of Data
Information Synthesis
Analysis

Visualization
Presentation of Data
Modeling, Simulation,
Image processing, Rendering

Los Alamos NATIONAL LABORATORY EST.1943

IBM

ASC NNSA

Industry direction towards C/GPU (Cell/B.E. is not alone)

AMD Direction toward Heterogeneous C/GPU (From Nov. 16, 2006)

The Fusion Vision

AMD Smarter Choice

Heterogeneous Processing:

AMD's Vision for CPU/GPU Integration

2008 Onward

Create the optimal computing experience for an increasingly mobile, graphics-centric world where digital media is becoming the norm

Deliver step-function improvements in microprocessor performance-per-watt-per-dollar over today's CPU-only architectures

Continue to scale x86 from palmtops to petaFLOPS by enabling new x86 computing paradigms, classes and form factors

October 2006 Unleashing the Processing Powerhouse

The Economic Tipping Point

AMD Smarter Choice

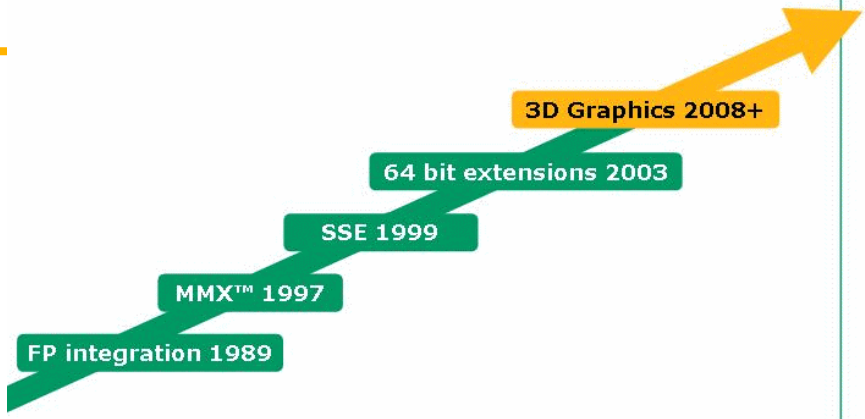

Now is the right time for CPU/GPU silicon-level integration

October 2006 Unleashing the Processing Powerhouse

Cell Like Heterogeneousness: GPU covers SPU functions

AMD's Fusion part Generation 1.0 C/GPU

The Data Efficiency Benefits of Silicon-Level Integration

AMD Smarter Choice

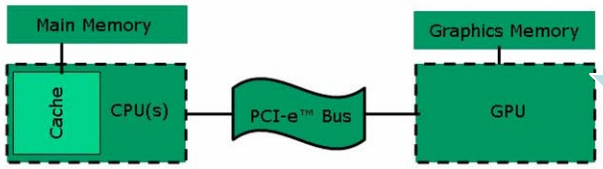

Maximum Power Usage

October 2006 Unleashing the Processing Powerhouse

EST.1943

Operated by the Los Alamos National Secur

The Data Efficiency Benefits of Silicon-Level Integration

AMD Smarter Choice

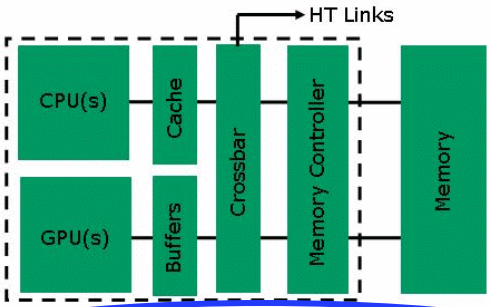

Expected Step-Function Improvement in Power/Performance

Unleashing the Processing Powerhouse

The CPU vs. GPU Debate Comes To An End!

October 2006 Unleashing the Processing Powerhouse

AMD Fusion

AMD Fusion is the codename for a future next-generation microprocessor design and the product of the merger between AMD and ATI, combining general processor execution as well as 3D geometry processing and other functions of today's GPUs into a single package. This technology is expected to debut in the timeframe of late 2008 or early 2009; as a successor of the latest microarchitecture, referred as "K8L". Four platforms focus on the four different aspects of usage

•General Purpose •Data Centric •Graphics Centric •Media Centric

Speed increase

There is to be an expected speed increase with the Fusion. Because the GPU and CPU will be on the same die, information transfer between the CPU and GPU/GPU memory will significantly increase since there will be no need for the information to travel on a bus as there is with current motherboards.

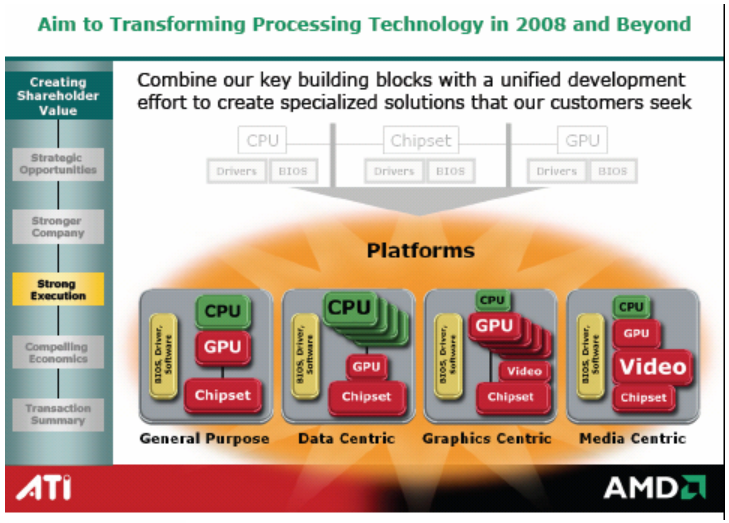

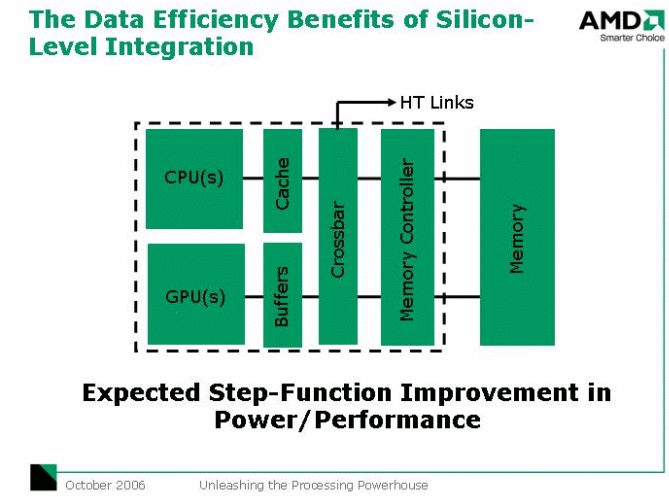

AMD's Direction C/GPU workloads (From Nov. 16, 2006)

Computational Processing is Evolving

- X86 computing continues to grow and expand its footprint but faces more challenges than ever:
  - Complex workloads and applications
  - Diverse operating environments
- New technologies are emerging to compliment X86 processing and meet the needs of heterogeneous processing environments
  - Stream Computing
  - Torrenza

Why Integrate CPU and GPU?

- CPUs and GPUs are both "PUs"
  - Processing Units
  - Both have similar demands on memory
    - Lots of Bandwidth
    - Lots of Capacity
    - Low Latency
  - Both take advantage of lots of fast transistors
- Many options to create great products using these two processing technologies coupled with a great semiconductor process and manufacturing story
- ..but you need a great "C" and a great "G" to make a compelling product

October 2006 Unleashing the Processing Powerhouse

AMD Stream Computing

- Allows graphics processors to accelerate complex computations in concert with the CPUs
  - Life sciences
  - Enterprise applications
  - Consumer applications

- AMD "Close to Metal"
  - New thin hardware interface that opens up stream architecture to developers
  - Low-level access to hardware to easily develop compilers, debuggers, math libraries, and application platforms.
- AMD Stream Processor
  - The world's first dedicated stream processor
  - Engineered specifically for compute-only systems such as workstations and servers

October 2006 Unleashing the Processing Powerhouse

From the Intel "Killer Future Apps Presentation"

Mining here refers to Rich Media Data

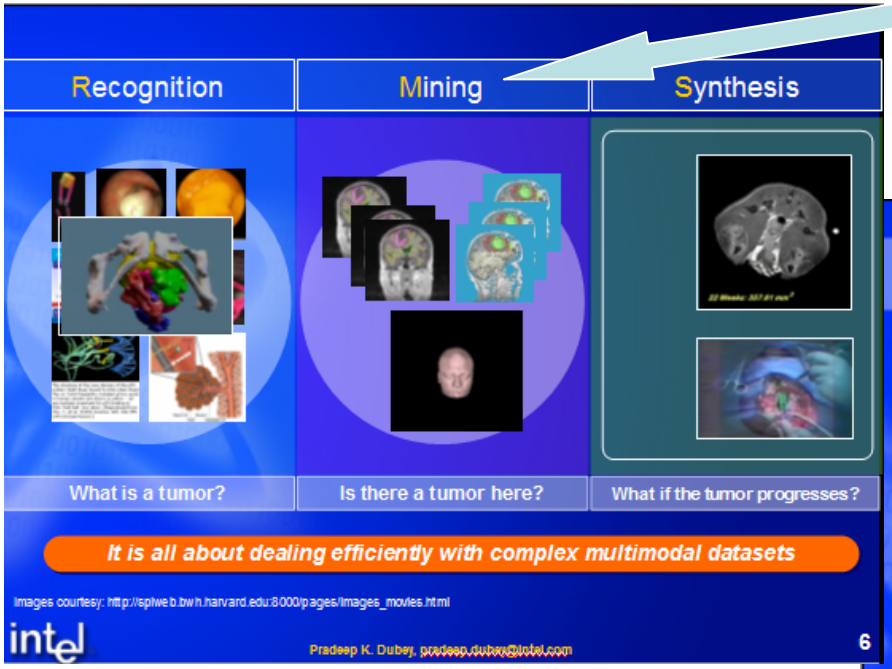

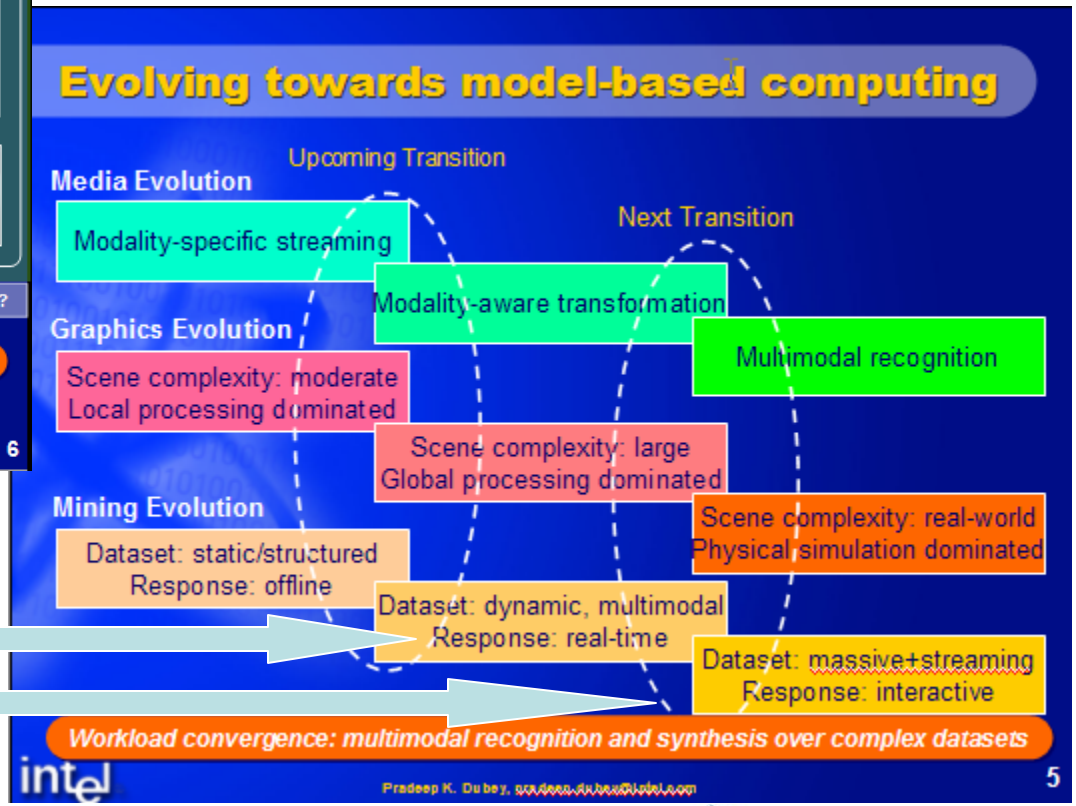

Real-Time,
Massive Data Sets,
Streaming

Intel Research on CPUs & GPUs and integration

Q&A With Jerry Bautista director of Intel's Microprocessor Research Lab

The Future Of CPUs & GPUs

September 2006 • Vol.6 Issue 9

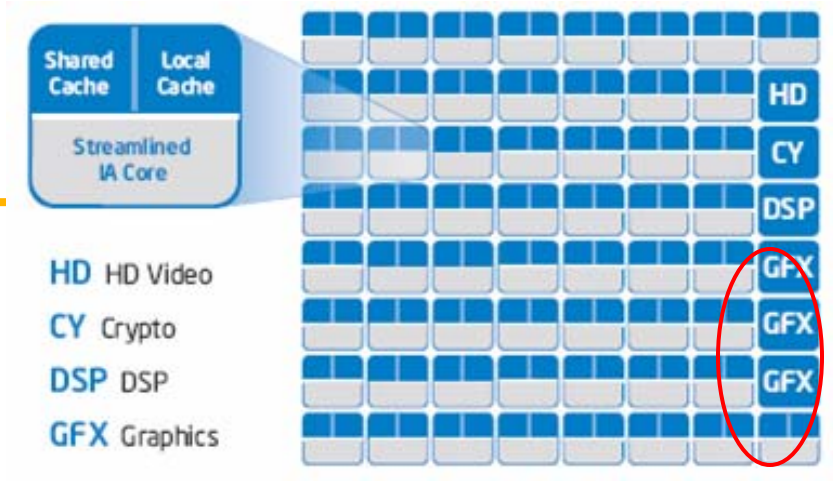

CPU: Jerry, do you believe that graphics will move back to the CPU?

JB: We see a trend. We watch the FLOPS (floating point operations), the watts, and dollars that go into the graphics cards and the computational physics on GPUs. They have been a growing part of the PC budget. We are aware of that. Some graphics computation is handled well on a graphics processor; **we can pull the graphics back on the CPU.**

CPU: Isn't there always going to be more than one processor in a system?

JB: The tera-scale computing project is where people miss the point. *It's not necessarily a homogeneous collection of cores.* The Cell microprocessor has one big core and eight synergistic processing elements; that is already a hybrid. We could have a general-purpose core with fixed function add-ons. You can do input/output acceleration, packet processing. There are a lot of doors. We know of execution doors with well-defined, discrete tasks. Why not build an engine to do those things? We know we would need lots of processors to do multidigital radio-signal processing. **Our version of tera-scale is a hybrid machine.**

CPU: What else do you use this computational ability for?

JB: Video searching with context. You can find a loved one in a lot of pictures at various lighting levels. These are sophisticated searches; like finding a face in an Interpol database. […] There is no end to the things we can do. The DARPA Grand Challenge (where they are trying to get a remotely controlled car to drive itself hundreds of miles) is a start. Having an autonomous vehicle is an example. Can you get them to navigate Manhattan traffic?

"Future CPU Architectures -- The Shift from Traditional Models Presentation" - 4Q'06 Douglas Carmean, Chief Architect, Intel's Visual Computing Group (VCG)

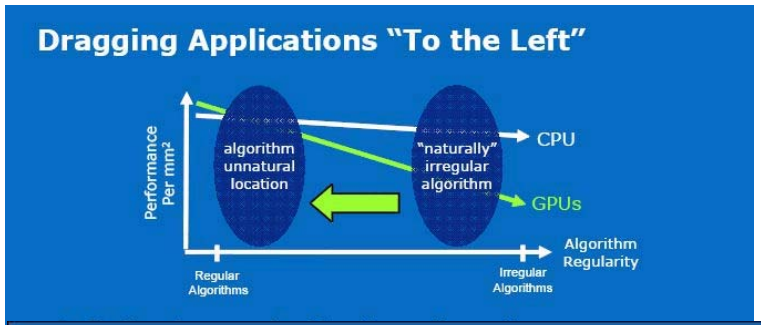

Project Larrabee — Intel has begun planning products based on a highly parallel, **IA-based programmable** architecture codenamed "Larrabee." It will be easily programmable using many existing software tools, and designed to scale to trillions of floating point operations per second (Teraflops) of performance. The Larrabee architecture will include enhancements to accelerate applications such as ***scientific computing, recognition, mining, synthesis, visualization, financial analytics and health applications***.

- Senior VP Pat Gelsinger - IDF Spring

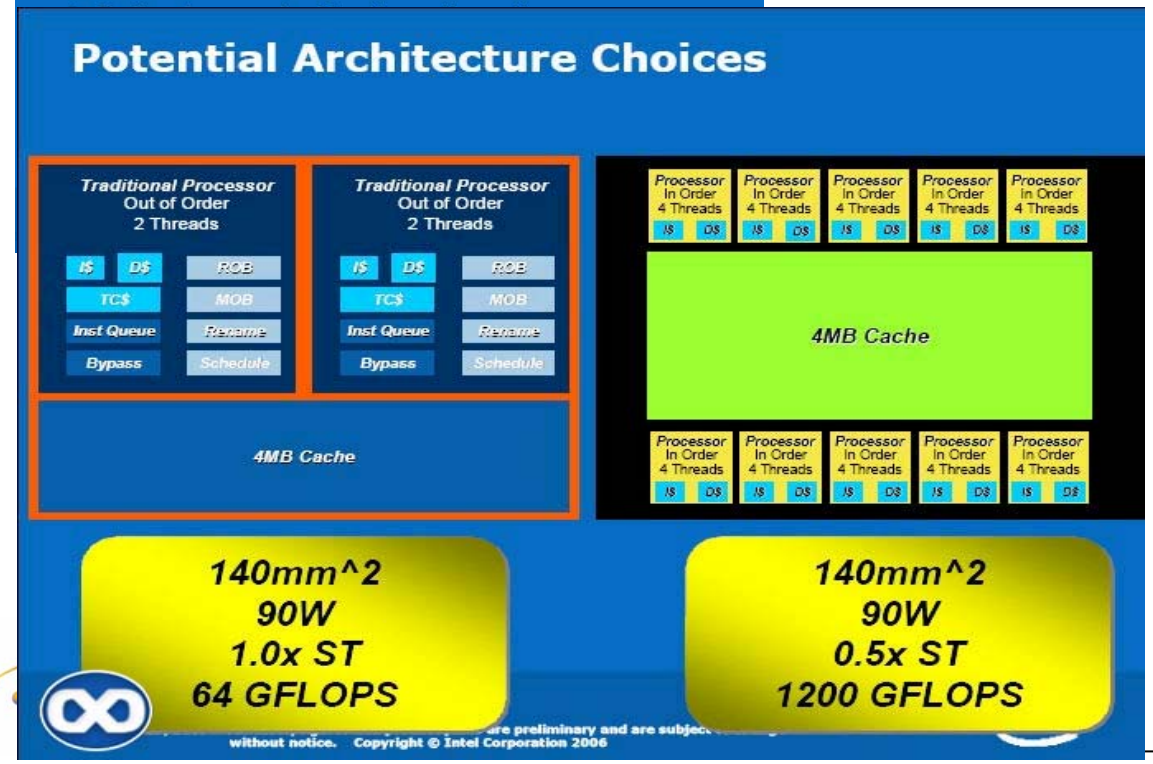

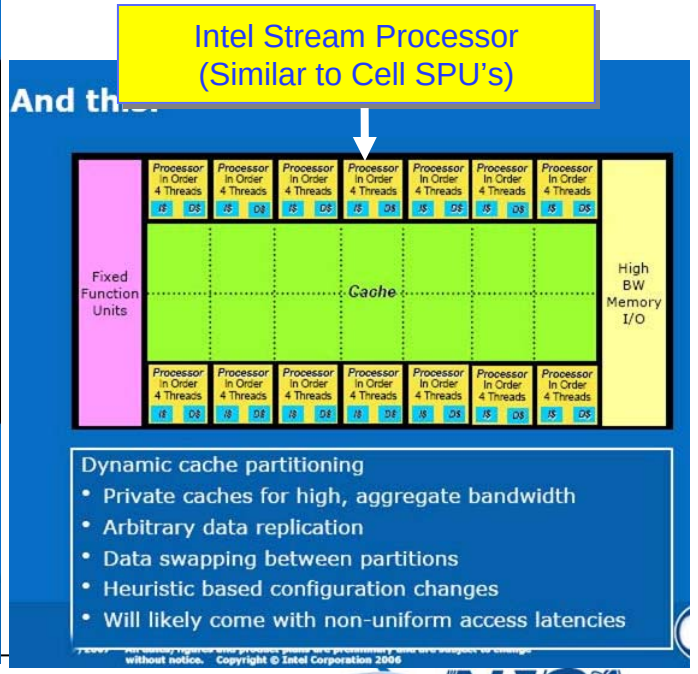

Evolution of nVidia GPUs toward Generalized Purpose

nVidia G70 - 2005
278 m Transistors

(15 months)

nVidia G80 - 2006
681 m Transistors

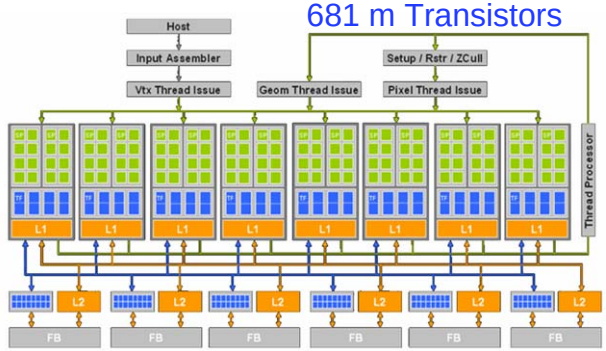

vs.

Cell - 2005
241 m Transistors

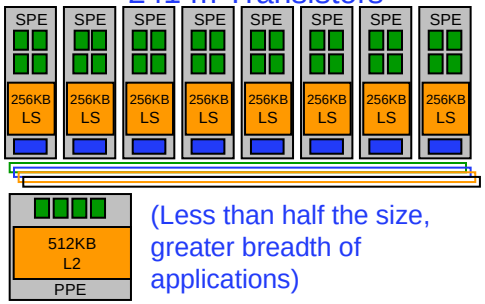

(Less than half the size, greater breadth of applications)

GPUs increasingly look like Cell/B.E.

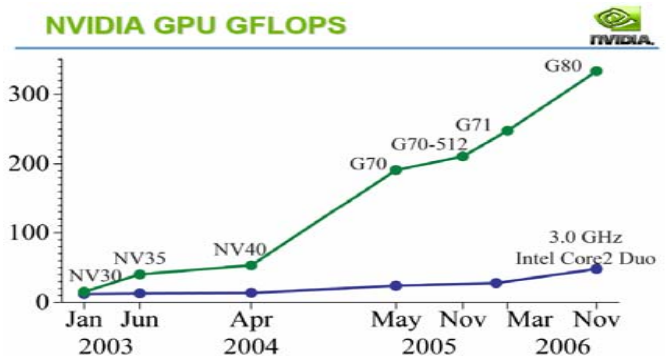

Chips Compared by area

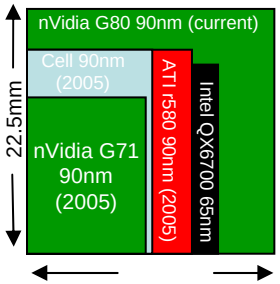

Chips Compared by type and transistor count

AMD Athlon 64 X2	CPU	154 m
Cell/B.E.	CGPU	241 m
Intel Core 2 Duo	CPU	291 m
Intel Pentium D 900	CPU	376 m
ATI X1950 XTX	GPU	384 m
Intel Core 2 Quad	CPU	582 m
NVIDIA G8800 GTX	GPU	681 m

Cell/B.E. competitive on flops/mm2 (and often more efficient)

Los Alamos NATIONAL LABORATORY EST.1943

IBM

ASC NNSA

RoadRunner vs. nVIDIA Tesla

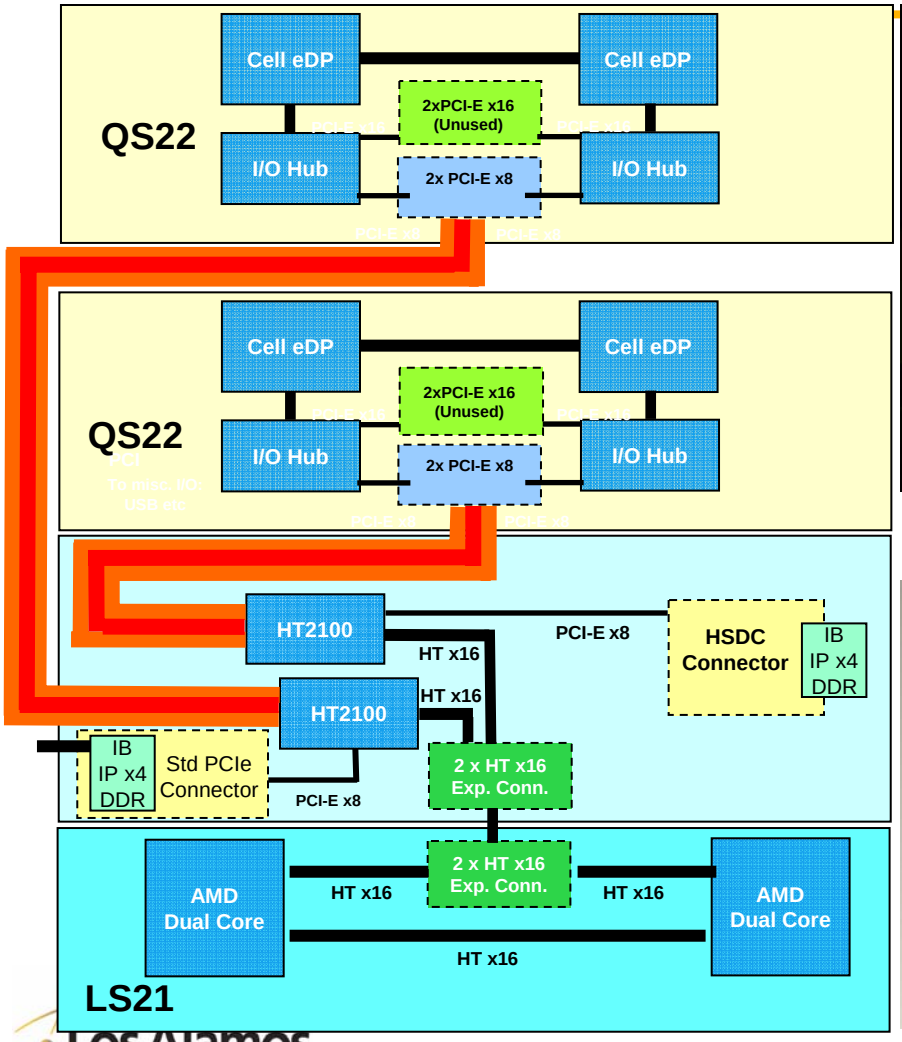

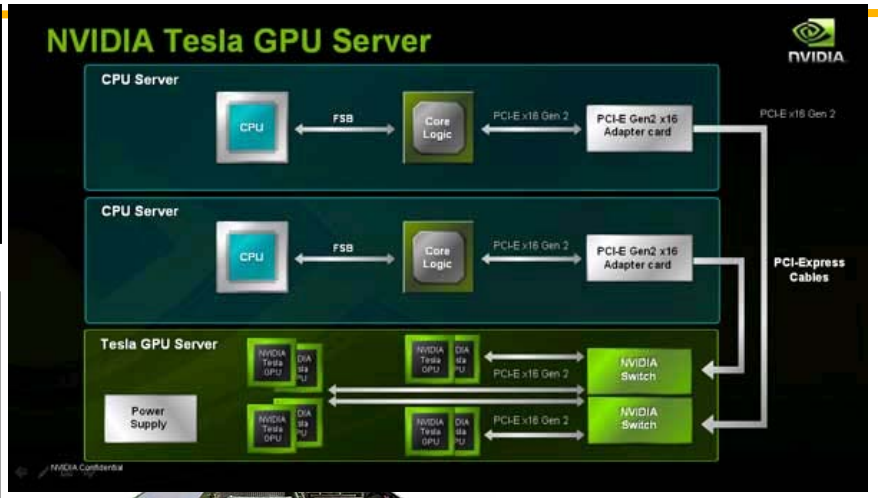

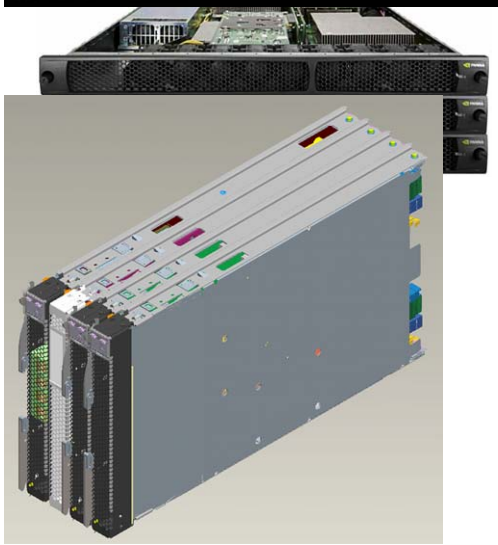

This diagram shows how the Tesla server configurations will work - one or more Tesla rack mount solutions working with one or more standard CPU-based servers that have PCI Express 2.0 x16 cards in them for communication with the Tesla servers. NVIDIA has said that to get the best efficiency out of these Tesla servers, there should be one CPU core per GPU core; that would mean a 4x4 server (16 cores) could support up to 16 GPUs (2-4 servers).

www.pcper.com

Roadrunner Advantages

- Significantly more memory per accelerator
  - Mitigates minimum offload size
- IEEE Double Precision Floating-Point
- Reliability & manageability
  - ECC on main memory and internal SRAM arrays
  - BladeCenter reliability and manageability
- Open architecture supported by open SDK
  - High level of participation from academia
  - No barriers to finding the best offload paradigms
- Large variety of Cell-based systems
  - Low barrier of entry for developers

Summary

Conclusions

PC Processors on the path towards hybrid architectures

… next step beyond multicore

... motivated by efficiency necessities and by a shift in workloads

Cell is not alone

… AMD and Intel both indicate hybrid model as the future

… Graphics processors evolving to become much more like SPEs

Cell is competitive on key metrics, peak flops/mm2 and flops/W vs. GPUs

… higher level of programmability than conventional GPUs

… higher performance densities than conventional CPUs

Special Notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Many of the features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special Notices (Cont.) -- Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: alphaWorks, BladeCenter, Blue Gene, ClusterProven, developerWorks, e business(logo), e(logo)business, e(logo)server, IBM, IBM(logo), ibm.com, IBM Business Partner (logo), IntelliStation, MediaStreamer, Micro Channel, NUMA-Q, PartnerWorld, PowerPC, PowerPC(logo), pSeries, TotalStorage, xSeries; Advanced Micro-Partitioning, eServer, Micro-Partitioning, NUMACenter, On Demand Business logo, OpenPower, POWER, Power Architecture, Power Everywhere, Power Family, Power PC, PowerPC Architecture, POWER5, POWER5+, POWER6, POWER6+, Redbooks, System p, System p5, System Storage, VideoCharger, Virtualization Engine.
A full list of U.S. trademarks owned by IBM may be found at: http://www.ibm.com/legal/copytrade.shtml.
Cell Broadband Engine, Cell/B.E., and PLAYSTATION are trademarks of Sony Computer Entertainment in the United States, other countries, or both.
Rambus is a registered trademark of Rambus, Inc.
XDR and FlexIO are trademarks of Rambus, Inc.
UNIX is a registered trademark in the United States, other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries or both.
Fedora is a trademark of Redhat, Inc.
Microsoft, Windows, Windows NT and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries or both.
Intel, Intel Xeon, Itanium and Pentium are trademarks or registered trademarks of Intel Corporation in the United States and/or other countries.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
AltiVec is a trademark of Freescale Semiconductor, Inc.
PCI-X and PCI Express are registered trademarks of PCI SIG.
InfiniBand™ is a trademark the InfiniBand® Trade Association
Other company, product and service names may be trademarks or service marks of others.